

CLAIMS

1. A method for determining where to store object replicas, the method comprising:

receiving an indication of a homeless replica of an object, wherein the object has a plurality of replicas including the homeless replica;

determining an initial placement for the homeless replica, wherein the initial placement is one of a plurality of devices in a system;

evaluating whether any replicas of the object can be swapped with one of a plurality of replicas of another object and not reduce a combined object availability of the two objects; and

swapping a replica of the object with the one of the plurality of replicas of the other object only if the swapping does not reduce the combined object availability of the two objects.

2. A method as recited in claim 1, wherein the replica of the object is stored on a first computing device, wherein the one of the plurality of replicas of the other object is stored on a second computing device, and wherein swapping the replicas comprises moving the replica of the object to the second computing device and moving the one of the plurality of replicas of the other object to the first computing device.

3. A method as recited in claim 1, wherein the method is implemented by multiple computing devices in a serverless distributed file system.

1 4. A method as recited in claim 1, wherein both the object and the other
2 object have a same amount of replicas.

3
4 5. A method as recited in claim 1, wherein the object comprises a file.

5
6 6. A method as recited in claim 1, wherein the object comprises a
7 plurality of files.

8
9 7. A method as recited in claim 1, wherein the object comprises a
10 portion of a file.

11
12 8. A method as recited in claim 1, wherein the object comprises a
13 directory.

14
15 9. A method as recited in claim 1, wherein the determining comprises
16 randomly selecting the initial placement for the homeless replica.

17
18 10. A method as recited in claim 1, wherein the evaluating comprises:
19 determining, for each possible swapping of replicas of the objects, an object
20 availability for each object; and
21 checking whether any of the possible swappings result in a combined object
22 availability of the two objects that is greater than the combined object availability
23 prior to the evaluating.

1 11. A method as recited in claim 1, wherein combined object
2 availability is not reduced if the availability of the object and the availability of the
3 other object are no further from one another than they were prior to the evaluating.

4
5 12. A method as recited in claim 1, wherein the swapping results in
6 increasing an availability of the object and decreasing an availability of the other
7 object.

8
9 13. A method as recited in claim 1, wherein the swapping results in
10 decreasing an availability of the object and increasing an availability of the other
11 object.

12
13 14. A method as recited in claim 1, implemented by a computing device
14 and further comprising selecting, prior to the evaluating, as the object one of a
15 plurality of objects that the computing device is responsible for maintaining.

16
17 15. A method as recited in claim 14, wherein the selecting comprises
18 randomly selecting one of the plurality of objects.

19
20 16. A method as recited in claim 14, wherein the selecting comprises
21 selecting the one of the plurality of objects that has the lowest object availability.

1 17. A method as recited in claim 1, wherein the other object is one of a
2 plurality of objects that a computing device is responsible for maintaining, and
3 wherein the other object is randomly selected from the plurality of objects.
4

5 18. A method as recited in claim 1, wherein the other object is one of a
6 plurality of objects that a computing device is responsible for maintaining, and
7 wherein the other object is the one of the plurality of objects that has the highest
8 object availability.
9

10 19. One or more computer readable media having stored thereon a
11 plurality of instructions that, when executed by one or more processors of a
12 computing device, causes the one or more processors to perform acts comprising:
13

14 working, in conjunction with one or more other processors of another
15 computing device, to determine whether a replica of a file managed by the
16 computing device and a replica of another file managed by the other computing
17 device can be swapped with one another to bring an availability of the file and an
18 availability of the other file closer; and

19 swapping the replica of the file and the replica of the other file only if the
20 swapping brings the availability of the file and the availability of the other file
21 closer.
22
23
24
25

1 **20.** One or more computer readable media as recited in claim 19,
2 wherein the swapping comprises communicating with a first device on which the
3 replica of the file is stored and a second device on which the replica of the other
4 file is stored to have the first device transfer the replica of the file to the second
5 device and delete the replica of the file on the first device, and to have the second
6 device transfer the replica of the other file to the first device and delete the replica
7 of the other file on the second device.

8
9 **21.** One or more computer readable media as recited in claim 19,
10 wherein the one or more processors are part of a device in a serverless distributed
11 file system.

12
13 **22.** One or more computer readable media as recited in claim 19,
14 wherein the file represents a plurality of files.

15
16 **23.** One or more computer readable media as recited in claim 19,
17 wherein the file represents a portion of another file.

18
19 **24.** One or more computer readable media as recited in claim 19,
20 wherein the file represents a folder.

21
22 **25.** One or more computer readable media as recited in claim 19,
23 wherein the plurality of instructions further cause the one or more processors to
24 perform acts comprising:

25 receiving an indication of a homeless replica of the file; and

1 determining an initial placement, on one of a plurality of additional
2 computing devices, for the homeless replica.

3
4 **26.** One or more computer readable media as recited in claim 25,
5 wherein determining the initial placement comprises randomly selecting the initial
6 placement for the homeless replica.

7
8 **27.** One or more computer readable media as recited in claim 25,
9 wherein the plurality of additional computing devices exclude any computing
10 device on which a replica of the file is already stored.

11
12 **28.** One or more computer readable media as recited in claim 19,
13 wherein the file is selected randomly from a plurality of files managed by the
14 computing device, and wherein the other file is selected randomly from another
15 plurality of files managed by the other computing device.

16
17 **29.** One or more computer readable media as recited in claim 19,
18 wherein the file is selected as the one of a plurality of files managed by the
19 computing device having a lowest file availability, and the other file is selected
20 randomly from another plurality of files managed by the other computing device.
21
22
23
24
25

1 **30.** One or more computer readable media as recited in claim 19,
2 wherein the file is selected as the one of a plurality of files managed by the
3 computing device having a lowest file availability, and the other file is selected as
4 the one of another plurality of files managed by the other computing device having
5 a highest file availability.
6

7 **31.** One or more computer readable media as recited in claim 19,
8 wherein the computing devices is part of a directory group that is collectively
9 responsible for managing a plurality of files including the file, and wherein the
10 other computing device is part of another directory group that is collectively
11 responsible for managing another plurality of files including the other file.
12

13 **32.** One or more computer readable media as recited in claim 19,
14 wherein the one or more processors of a computing device and the one or more
15 other processors of another computing device are the same processors of the same
16 computing device.
17

18 **33.** A serverless distributed file system comprising:
19 a first plurality of computing devices storing files;
20 a second plurality of computing devices managing storage of the files;
21 wherein a first computing device of the second plurality of computing
22 devices selects a file for which it manages storage and communicates with a
23 second computing device of the second plurality of computing devices to
24 determine whether a replica of the file and a replica of another file for which the
25

1 second computing device manages storage can be swapped in order to improve a
2 combined file availability; and

3 if the replicas can be swapped to improve the combined file availability,
4 then instructing the one of the first plurality of computing devices on which the
5 replica of the file is stored to transfer the replica of the file to the one of the first
6 plurality of computing devices on which the replica of the other file is stored, and
7 instructing the one of the first plurality of computing devices on which the replica
8 of the other file is stored to transfer the replica of the other file to the one of the
9 first plurality of computing devices on which the replica of the file is stored.

10
11 **34.** A serverless distributed file system as recited in claim 33, wherein
12 the second plurality of computing devices comprise one or more directory groups.

13
14 **35.** A serverless distributed file system as recited in claim 33, wherein
15 the second plurality of computing devices further receives an indication of a
16 homeless replica of the file, and randomly selects, as a computing device on which
17 to store the homeless replica, one of the first plurality of computing devices on
18 which no other replica of the file is already stored.

19
20 **36.** One or more computer readable media having stored thereon a
21 plurality of instructions that, when executed by one or more processors of a
22 computing device, causes the one or more processors to:

23 initially place replicas of a file on different ones of a plurality of devices
24 using a first process; and
25

1 subsequently improve the placement of replicas of a plurality of files by
2 evaluating whether any replicas of a first file can be swapped with any replicas of
3 a second file without a reduction in the combined file availability of the first and
4 second files, and swapping a replica of the first file with a replica of the second
5 file if the swapping results in no reduction in the combined file availability of the
6 first and second files.

7
8 **37.** One or more computer readable media as recited in claim 36,
9 wherein the first process comprises random selection.

10
11 **38.** One or more computer readable media as recited in claim 36,
12 wherein swapping the replica of the first file with the replica of the second file
13 comprises communicating with a first device on which the replica of the first file
14 is stored and a second device on which the replica of the second file is stored to
15 have the first device transfer the replica of the first file to the second device and
16 delete the replica of the first file on the first device, and to have the second device
17 transfer the replica of the second file to the first device and delete the replica of the
18 second file on the second device.

19
20 **39.** One or more computer readable media as recited in claim 36,
21 wherein the first file is selected randomly from a first plurality of files, and
22 wherein the second file is selected randomly from a second plurality of files.
23
24
25

1 **40.** One or more computer readable media as recited in claim 36,
2 wherein the first file is selected as the one of a first plurality of files having a
3 lowest file availability, and the second file is selected randomly from a second
4 plurality of files.

5
6 **41.** One or more computer readable media as recited in claim 36,
7 wherein the first file is selected as the one of a first plurality of files managed
8 having a lowest file availability, and the second file is selected as the one of a
9 second plurality of files having a highest file availability.

10
11 **42.** A method, implemented in a directory group, the method
12 comprising:

13 selecting another directory group to participate with in a replica placement
14 improvement process;

15 selecting a file maintained by the directory group;

16 determining whether exchanging a replica of the file with a replica of
17 another file maintained by the other directory group will increase a combined file
18 availability of the files; and

19 having the replica of the file and the replica of the other file exchanged if
20 exchanging the replicas will increase the combined file availability of the files.

21
22 **43.** A method as recited in claim 42, further comprising:

23 receiving, at the directory group, an indication of a homeless replica of the
24 file; and
25

1 selecting, as a computing device on which to store the homeless replica,
2 one of a plurality of computing devices on which no other replica of the file is
3 already stored.

4
5 **44.** A method as recited in claim 42, wherein the selecting comprises
6 randomly selecting.

7
8 **45.** A method as recited in claim 42, wherein having the replica of the
9 file and the replica of the other file exchanged comprises communicating with a
10 first device on which the replica of the file is stored and a second device on which
11 the replica of the other file is stored to have the first device transfer the replica of
12 the file to the second device and delete the replica of the file on the first device,
13 and to have the second device transfer the replica of the other file to the first
14 device and delete the replica of the other file on the second device.

15
16 **46.** A method as recited in claim 42, wherein the file is selected
17 randomly from a plurality of files managed by the directory group, and wherein
18 the other file is selected randomly from another plurality of files managed by the
19 other directory group.

20
21 **47.** A method as recited in claim 42, wherein the file is selected as the
22 one of a plurality of files managed by the directory group having a lowest file
23 availability, and the other file is selected randomly from another plurality of files
24 managed by the other directory group.
25

1 **48.** A method as recited in claim 42, wherein the file is selected as the
2 one of a plurality of files managed by the directory group having a lowest file
3 availability, and the other file is selected as the one of another plurality of files
4 managed by the other directory group having a highest file availability.
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25